arXiv:1409.6873v3 [cs.LO] 22 Jul 2015

Probabilistic Thread Algebra

J.A. Bergstra and C.A. Middelburg

Informatics Institute, Faculty of Science, University of Amsterdam,
Science Park 904, 1098 XH Amsterdam, the Netherlands
J.A.Bergstra@uva.nl,C.A.Middelburg@uva.nl

Abstract. We add probabilistic features to basic thread algebra and its extensions with thread-service interaction and strategic interleaving. Here, threads represent the behaviours produced by instruction sequences under execution and services represent the behaviours exhibited by the components of execution environments of instruction sequences. In a paper concerned with probabilistic instruction sequences, we proposed several kinds of probabilistic instructions and gave an informal explanation for each of them. The probabilistic features added to the extension of basic thread algebra with thread-service interaction make it possible to give a formal explanation in terms of non-probabilistic instructions and probabilistic services. The probabilistic features added to the extensions of basic thread algebra with strategic interleaving make it possible to cover strategies corresponding to probabilistic scheduling algorithms.

Keywords: basic thread algebra, probabilistic thread, probabilistic service, probabilistic interleaving strategy, probabilistic instruction.

1998 ACM Computing Classification: D.3.3, D.4.1, F.1.1, F.1.2.

1 Introduction

In [6], an approach to the semantics of programming languages was presented which is based on the perspective that a program is in essence an instruction sequence. The groundwork for the approach is formed by PGA (ProGram Algebra), an algebraic theory of single-pass instruction sequences, and BTA (Basic Thread Algebra), an algebraic theory of mathematical objects that represent the behaviours produced by instruction sequences under execution (for a comprehensive introduction to these algebraic theories, see [15]). To increase the applicability of the approach, BTA was extended with thread-service interaction in [16]. In the setting of BTA and its extension with thread-service interaction, threads are mathematical objects that represent the behaviours produced by instruction sequences under execution and services are mathematical objects that represent the behaviours exhibited by components of execution environments of instruction sequences.

As a continuation of the work presented in [6,16], (a) the notion of an instruction sequence was subjected to systematic and precise analysis using the

groundwork laid earlier, (b) various issues, including issues relating to computability and complexity of computational problems, efficiency of algorithms, and verification of programs, were rigorously investigated thinking in terms of instruction sequences (for a comprehensive survey of a large part of the work referred to under (a) and (b), see [15]), and (c) the form of interleaving concurrency that is relevant to the behaviours of multi-threaded programs under execution, called strategic interleaving in the setting of BTA, was rigorously investigated by means of extensions of BTA (see e.g. [8,9,10]).

In the course of the work referred to above under (b), we ran into the problem that BTA and its extension with thread-service interaction do not allow issues relating to probabilistic computation to be investigated thinking in terms of instruction sequences. In the course of the work referred to above under (c), we ran into the problem that BTA also does not allow probabilistic strategic interleaving to be investigated by means of extensions of BTA. This paper concerns the addition of features to BTA and its extensions with thread-service interaction and strategic interleaving that will take away these limitations.

We consider it important to take probabilistic computation into account in future investigations. The primary reasons for this are the following: (a) the existence of probabilistic algorithms that are highly efficient, possibly at the cost of a probability of correctness less than one (e.g. primality testing, see [29]); (b) the existence of probabilistic algorithms for which no deterministic counterparts exist (e.g. symmetry breaking, see [26]); (c) the gradually created evidence for the hypothesis that it is relevant for a diversity of issues in computer science and engineering to think in terms of instruction sequences. This constitutes the basis of our motivation for the work presented in this paper.

In [12], we gave an enumeration of kinds of probabilistic instructions that were chosen on the basis of direct intuitions and therefore not necessarily the best kinds in any sense. We only gave an informal explanation for each of the enumerated kinds because we considered it premature at the time to add probabilistic features to BTA that would make it possible to give a formal explanation. We were doubtful whether the ad hoc addition of features to BTA was the right way to go.

Later, we have found that the ramification of semantic options with the addition of probabilistic features to BTA is well surveyable because of (a) the limitation of the scope to behaviours produced by instruction sequences under execution and (b) the semantic constraints brought about by the informal explanations of the kinds of probabilistic instructions enumerated in [12] and the desired elimination property of all but one kind. In the case of a general process algebra, such as ACP [3], CCS [27] or CSP [24], the ramification becomes much more complex, particularly because a limitation of the scope to behaviours of a special kind is lacking. In this paper, we add probabilistic features to BTA and an extension of BTA with thread-service interaction.

The probabilistic features added to the extension of BTA with thread-service interaction make it possible to give a formal explanation for each of the kinds of probabilistic instructions enumerated in [12] in terms of non-probabilistic in-

structions and probabilistic services. To demonstrate this, we add the kind of probabilistic instructions that cannot be eliminated to PGLB (ProGramming Language B), a program notation rooted in PGA and close to existing assembly languages, and give a formal definition of the behaviours produced by the instruction sequences from the resulting program notation. We opted for PGLB because in the past it has proved itself suitable for the investigation of various issues. The added kind of probabilistic instructions allow probabilistic choices to be made during the execution of instruction sequences.

In [8] and subsequent papers, we extended BTA with kinds of interleaving where interleaving takes place according to some deterministic interleaving strategy. Interleaving strategies are abstractions of scheduling algorithms. Interleaving according to an interleaving strategy differs from arbitrary interleaving, but it is what really happens in the case of multi-threading as found in programming languages such as Java [21] and C# [23]. The extension of BTA with a probabilistic feature does not only allow of probabilistic services, but also allows of probabilistic interleaving strategies. In this paper, we also generalize the extensions of BTA with specific kinds of deterministic strategic interleaving to an extension for a large class of kinds of deterministic and probabilistic strategic interleaving. Thus, strategies corresponding to probabilistic scheduling algorithms such as the lottery scheduling algorithm [33] are covered.

The main results of this paper are probabilistic versions of BTA and its extensions with thread-service interaction and strategic interleaving which pave the way for (a) investigation of issues related to probabilistic computation thinking in terms of instruction sequences and (b) investigation of probabilistic interleaving strategies.

In this paper, we take functions whose range is the carrier of a signed cancellation meadow as probability measures. In [18], meadows are proposed as alternatives for fields with a purely equational axiomatization. A meadow is a commutative ring with a multiplicative identity element and a total multiplicative inverse operation satisfying two equations which imply that the multiplicative inverse of zero is zero. A cancellation meadow is a field whose multiplicative inverse operation is made total by imposing that the multiplicative inverse of zero is zero, and a signed cancellation meadow is a cancellation meadow expanded with a signum operation. In [17], Kolmogorov's probability axioms for finitely additive probability spaces are rephrased for the case where probability measures are functions whose range is the carrier of a signed cancellation meadow.

This paper is organized as follows. First, we review signed cancellation meadows (Section 2). Next, we add probabilistic features to BTA and an extension of BTA with thread-service interaction (Sections 3 and 4). Then, we add a kind of probabilistic instructions to PGLB (Section 5). Following this, we add probabilistic features to the extensions of BTA with strategic interleaving (Section 6). Finally, we make some concluding remarks (Section 7).

It should be mentioned that BTA is introduced in [6] under the name BPPA (Basic Polarized Process Algebra) and services are called state machines in [16].

2 Signed Cancellation Meadows

We will take functions whose range is the carrier of a signed cancellation meadow as probability measures. Therefore, we review signed cancellation meadows in this section.

In [18], meadows are proposed as alternatives for fields with a purely equational axiomatization. A meadow is a commutative ring with a multiplicative identity element and a total multiplicative inverse operation satisfying two equations which imply that the multiplicative inverse of zero is zero. Thus, all meadows are total algebras and the class of all meadows is a variety. At the basis of meadows lies the decision to make the multiplicative inverse operation total by imposing that the multiplicative inverse of zero is zero. All fields in which the multiplicative inverse of zero is zero, called zero-totalized fields, are meadows, but not conversely.

A cancellation meadow is a meadow that satisfies the *cancellation axiom* $x \neq 0 \wedge x \cdot y = x \cdot z \Rightarrow y = z$. The zero-totalized fields are exactly the cancellation meadows that satisfy in addition the *separation axiom* $0 \neq 1$. A paradigmatic example of cancellation meadows is the field of rational numbers with the multiplicative inverse operation made total by imposing that the multiplicative inverse of zero is zero (see e.g. [18]). An example of a meadow that is not a zero-totalized field is the initial algebra of the equational axiomatization of meadows (see e.g. [5]).

A signed cancellation meadow is a cancellation meadow expanded with a signum operation. The usefulness of the signum operation lies in the fact that the predicates $<$ and $\leq$ can be defined using this operation (see below).

The signature of signed cancellation meadows consists of the following constants and operators: the constants 0 and 1, the binary *addition* operator $+$, the binary *multiplication* operator $\cdot$, the unary *additive inverse* operator $-$, the unary *multiplicative inverse* operator $^{-1}$, and the unary *signum* operator s.

Terms are built as usual. We use infix notation for the binary operators $+$ and $\cdot$, prefix notation for the unary operator $-$, and postfix notation for the unary operator $^{-1}$. We use the usual precedence convention to reduce the need for parentheses. We introduce subtraction and division as abbreviations: $t - t'$ abbreviates $t + (-t')$ and t/t' abbreviates $t \cdot (t'^{-1})$.

The constants and operators from the signature of signed cancellation meadows are adopted from rational arithmetic, which gives an appropriate intuition about these constants and operators.

Signed cancellation meadows are axiomatized by the equations in Tables 1 and 2 and the above-mentioned cancellation axiom. The axioms for the signum operator stem from [4].

The predicates $<$ and $\leq$ are defined in signed cancellation meadows as follows: $x < y \Leftrightarrow \mathsf{s}(y-x) = 1$ and $x \leq y \Leftrightarrow \mathsf{s}(\mathsf{s}(y-x)+1) = 1$. Because $\mathsf{s}(\mathsf{s}(y-x)+1) \neq -1$, we have $0 \leq x \leq 1 \Leftrightarrow \mathsf{s}(\mathsf{s}(x)+1) \cdot \mathsf{s}(\mathsf{s}(1-x)+1) = 1$. We will use this equivalence below to describe the set of probabilities.

Table 1. Axioms of a meadow

$(x+y)+z = x+(y+z)$	$(x \cdot y) \cdot z = x \cdot (y \cdot z)$	$(x^{-1})^{-1} = x$
$x+y = y+x$	$x \cdot y = y \cdot x$	$x \cdot (x \cdot x^{-1}) = x$
$x+0 = x$	$x \cdot 1 = x$	
$x+(-x) = 0$	$x \cdot (y+z) = x \cdot y + x \cdot z$	

Table 2. Additional axioms for the signum operator

$\mathsf{s}(x/x) = x/x$	$\mathsf{s}(x^{-1}) = \mathsf{s}(x)$
$\mathsf{s}(1-x/x) = 1-x/x$	$\mathsf{s}(x \cdot y) = \mathsf{s}(x) \cdot \mathsf{s}(y)$
$\mathsf{s}(-1) = -1$	$(1 - \frac{\mathsf{s}(x)-\mathsf{s}(y)}{\mathsf{s}(x)-\mathsf{s}(y)}) \cdot (\mathsf{s}(x+y) - \mathsf{s}(x)) = 0$

3 Probabilistic Basic Thread Algebra

In this section, we introduce prBTA (probabilistic Basic Thread Algebra), a probabilistic version of BTA. The objects considered in BTA are called threads. In BTA, a thread represents a behaviour which consists of performing actions in a deterministic sequential fashion. Upon each action performed, a reply from an execution environment determines how the thread proceeds. The possible replies are the values t and f. In prBTA, a thread represents a behaviour which consists of performing actions in a probabilistic sequential fashion. That is, performing actions may alternate with making internal choices according to discrete probability distributions.

In the sequel, it is assumed that a fixed but arbitrary signed cancellation meadow $\mathfrak{M}$ has been given. We denote the carrier of $\mathfrak{M}$ by $\mathfrak{M}$ as well, and we denote the interpretations of the constants and operators in $\mathfrak{M}$ by the constants and operators themselves. We write $\mathcal{P}$ for the set $\{\pi \in \mathfrak{M} \mid \mathsf{s}(\mathsf{s}(\pi)+1) \cdot \mathsf{s}(\mathsf{s}(1-\pi)+1) = 1\}$ of probabilities.

In prBTA, it is moreover assumed that a fixed but arbitrary set $\mathcal{A}$ of *basic actions*, with $\mathsf{tau} \notin \mathcal{A}$, has been given. In addition, there is the special action tau. Performing tau, which is considered performing an internal action, will always lead to the reply t. We write $\mathcal{A}_{\mathsf{tau}}$ for $\mathcal{A} \cup \{\mathsf{tau}\}$ and refer to the members of $\mathcal{A}_{\mathsf{tau}}$ as basic actions.

The algebraic theory prBTA has one sort: the sort $\mathbf{T}$ of *threads*. We make this sort explicit to anticipate the need for many-sortedness later on. To build terms of sort $\mathbf{T}$, prBTA has the following constants and operators:

- the *inaction* constant $\mathsf{D} : \to \mathbf{T}$;[1]
- the *termination* constant $\mathsf{S} : \to \mathbf{T}$;
- for each $a \in \mathcal{A}_{\mathsf{tau}}$, the binary *postconditional composition* operator $_ \trianglelefteq a \trianglerighteq _ : \mathbf{T} \times \mathbf{T} \to \mathbf{T}$;

[1] In earlier work, the inaction constant is sometimes called the deadlock constant.

Table 3. Axioms of prBTA

$x \trianglelefteq \mathsf{tau} \trianglerighteq y = x \trianglelefteq \mathsf{tau} \trianglerighteq x$	T1
$x +_\pi y = y +_{1-\pi} x$	prA1
$x +_\pi (y +_\rho z) = (x +_{\frac{\pi}{\pi+\rho-\pi\cdot\rho}} y) +_{\pi+\rho-\pi\cdot\rho} z$	prA2
$x +_\pi x = x$	prA3
$x +_1 y = x$	prA4

– for each $\pi \in \mathcal{P}$, the binary *probabilistic composition* operator $_ +_\pi _ : \mathbf{T} \times \mathbf{T} \to \mathbf{T}$.

Terms of sort $\mathbf{T}$ are built as usual in the one-sorted case. We assume that there are infinitely many variables of sort $\mathbf{T}$, including x, y, z. We use infix notation for postconditional composition and probabilistic composition. We introduce *basic action prefixing* as an abbreviation: $a \circ t$, where t is a prBTA term, abbreviates $t \trianglelefteq a \trianglerighteq t$. We identify expressions of the form $a \circ t$ with the prBTA terms they stand for.

The thread denoted by a closed term of the form $t \trianglelefteq a \trianglerighteq t'$ will first perform a, and then proceed as the thread denoted by t if the reply from the execution environment is t and proceed as the thread denoted by t' if the reply from the execution environment is f. The thread denoted by a closed term of the form $t +_\pi t'$ will behave like the thread denoted by t with probability π and like the thread denoted by t' with probability $1 - \pi$. The thread denoted by S will do no more than terminate and the thread denoted by D will become inactive. A thread becomes inactive if no more basic actions are performed, but it does not terminate.

The inaction constant, the termination constant and the postconditional composition operators are adopted from BTA. Counterparts of the probabilistic composition operators are found in most probabilistic process algebras that offer probabilistic choices of the generative variety (see e.g. [2]).

The axioms of prBTA are given in Table 3. In this table, π and ρ stand for arbitrary probabilities from $\mathcal{P}$. Axiom T1 reflects that performing tau will always lead to the reply t and axioms prA1–prA4 express that probabilistic composition provides probabilistic choices of the generative variety (see [20]). From prA1 and prA4, we can derive both $x +_0 (y +_0 z) = z$ and $(x +_0 y) +_0 z = z$, and hence also $x +_0 (y +_0 z) = (x +_0 y) +_0 z$. This last equation can be immediately derived from prA2 as well because in meadows $0/0 = 0$.

Axiom T1 is adopted from BTA. Counterparts of axioms prA1–prA3 are found in most probabilistic process algebras that offer probabilistic choices of the generative variety (see e.g. [2]). However, in the process algebras concerned the probabilities 0 and 1 are excluded in probabilistic choices to prevent division by zero. Owing to this exclusion, axiom prA4 is lacking in these process algebras.

Each closed prBTA term denotes a finite thread, i.e. a thread with a finite upper bound to the number of basic actions that it can perform. Infinite threads,

Table 4. Axioms for the guarded recursion constants

$\langle X\|E\rangle = \langle t_X\|E\rangle$	if $X = t_X \in E$	RDP
$E \Rightarrow X = \langle X\|E\rangle$	if $X \in \mathrm{V}(E)$	RSP

i.e. threads without a finite upper bound to the number of basic actions that it can perform, can be described by guarded recursion. A *guarded recursive specification* over prBTA is a set of recursion equations $E = \{X = t_X \mid X \in V\}$, where V is a set of variables of sort $\mathbf{T}$ and each t_X is a prBTA term in which only variables from V occur and each occurrence of a variable in t_X is in a subterm of the form $t \trianglelefteq a \trianglerighteq t'$. We write $\mathrm{V}(E)$ for the set of all variables that occur on the left-hand side of an equation in E.

We are only interested in models of prBTA in which guarded recursive specifications have unique solutions. A model of prBTA in which guarded recursive specifications have unique solutions is the projective limit model of prBTA. This model is constructed along the same line as the projective limit model of BTA presented in [15]. It is based on the view that two threads are identical if their approximations up to any finite depth are identical. The approximation up to depth n of a thread is obtained by cutting it off after it has performed n actions if it has not yet terminated or become inactive.

We confine ourselves to the projective limit model of prBTA, which has an initial model of prBTA as a submodel, for the interpretation of prBTA terms. An outline of this model is given in Appendix A.1. In the sequel, we use the term *probabilistic thread* or simply *thread* for the elements of the carrier of the model. Regular threads, i.e. finite or infinite threads that can only be in a finite number of states, can be defined by means of a finite guarded recursive specification.

We extend prBTA with guarded recursion by adding constants for solutions of guarded recursive specifications and axioms concerning these additional constants. For each guarded recursive specification E and each $X \in \mathrm{V}(E)$, we add a constant standing for the unique solution of E for X to the constants of prBTA. The constant standing for the unique solution of E for X is denoted by $\langle X|E\rangle$. Moreover, we use the following notation. Let t be a prBTA term and E be a guarded recursive specification. Then we write $\langle t|E\rangle$ for t with, for all $X \in \mathrm{V}(E)$, all occurrences of X in t replaced by $\langle X|E\rangle$. We add the axioms for guarded recursion given in Table 4 to the axioms of prBTA. In this table, X, t_X and E stand for an arbitrary variable of sort $\mathbf{T}$, an arbitrary prBTA term and an arbitrary guarded recursive specification, respectively. Side conditions are added to restrict the variables, terms and guarded recursive specifications for which X, t_X and E stand.

The additional axioms for guarded recursion are known as the recursive definition principle (RDP) and the recursive specification principle (RSP). The equations $\langle X|E\rangle = \langle t_X|E\rangle$ for a fixed E express that the constants $\langle X|E\rangle$ make up a solution of E. The conditional equations $E \Rightarrow X = \langle X|E\rangle$ express that this solution is the only one.

In Section 6, we will use the notation $\sum_{i=k}^{n}[\pi_i]\,t_i$ with $1 \leq k \leq n$ and $\sum_{i=k}^{n}\pi_i = 1$ for right-nested probabilistic composition. The term $\sum_{i=k}^{n}[\pi_i]\,t_i$ with $1 \leq k \leq n$ is defined by induction on $n-k$ as follows:

$$\begin{aligned}
&\textstyle\sum_{i=k}^{n}[\pi_i]\,t_i = t_k && \text{if } k = n\ ,\\
&\textstyle\sum_{i=k}^{n}[\pi_i]\,t_i = t_k +_{\pi_k} (\sum_{i=k+1}^{n}[\tfrac{\pi_i}{1-\pi_k}]\,t_i) && \text{if } k < n\ .
\end{aligned}$$

The thread denoted by $\sum_{i=k}^{n}[\pi_i]\,t_i$ will behave like the thread denoted by t_k with probability π_k and ... and like the thread denoted by t_n with probability π_n.

4 Interaction of Threads with Services

Services are objects that represent the behaviours exhibited by components of execution environments of instruction sequences at a high level of abstraction. A service is able to process certain methods. The processing of a method may involve a change of the service. At completion of the processing of a method, the service produces a reply value. Execution environments are considered to provide a family of uniquely-named services. A thread may interact with the named services from the service family provided by an execution environment. That is, a thread may perform a basic action for the purpose of requesting a named service to process a method and to return a reply value at completion of the processing of the method. In this section, we extend prBTA with services, service families, a composition operator for service families, an operator that is concerned with this kind of interaction, and a general operator for abstraction from the internal action tau.

In SFA, the algebraic theory of service families introduced below, it is assumed that a fixed but arbitrary set $\mathcal{M}$ of *methods* has been given. Moreover, the following is assumed with respect to services:

- a signature $\Sigma_{\boldsymbol{\mathcal{S}}}$ has been given that includes the following sorts:
  - the sort $\mathbf{S}$ of *services*;
  - the sort $\mathbf{B}$ of *Boolean values*;

  and the following constants and operators:
  - the *empty service* constant $\delta :\to \mathbf{S}$;
  - the *reply* constants $\mathsf{t}, \mathsf{f} :\to \mathbf{B}$;
  - for each $m \in \mathcal{M}$, the *derived service* operator $\frac{\partial}{\partial m} : \mathbf{S} \to \mathbf{S}$;
  - for each $m \in \mathcal{M}$ and $\pi \in \mathcal{P}$, the *service reply* operator $\varrho_m^\pi : \mathbf{S} \to \mathbf{B}$;
- a minimal $\Sigma_{\boldsymbol{\mathcal{S}}}$-algebra $\boldsymbol{\mathcal{S}}$ has been given in which the following holds:
  - $\mathsf{t} \neq \mathsf{f}$;
  - $\bigwedge_{m\in\mathcal{M}} (\frac{\partial}{\partial m}(s) = \delta \Leftrightarrow \bigwedge_{\pi\in\mathcal{P}} \varrho_m^\pi(s) = \mathsf{f})$;
  - $\bigwedge_{\pi,\rho\in\mathcal{P}} (\varrho_m^\pi(s) = \mathsf{t} \wedge \varrho_m^\rho(s) = \mathsf{t} \Rightarrow \pi = \rho)$.

The intuition concerning $\frac{\partial}{\partial m}$ and ϱ_m^π is that on a request to service s to process method m:

- if $\varrho_m^\pi(s) = \mathsf{t}$, s processes m, produces the reply t with probability π and the reply f with probability $1-\pi$, and then proceeds as $\frac{\partial}{\partial m}(s)$;

Table 5. Axioms of SFA

$u \oplus \emptyset = u$	SFC1	$\partial_F(\emptyset) = \emptyset$	SFE1
$u \oplus v = v \oplus u$	SFC2	$\partial_F(f.s) = \emptyset$ if $f \in F$	SFE2
$(u \oplus v) \oplus w = u \oplus (v \oplus w)$	SFC3	$\partial_F(f.s) = f.s$ if $f \notin F$	SFE3
$f.s \oplus f.s' = f.\delta$	SFC4	$\partial_F(u \oplus v) = \partial_F(u) \oplus \partial_F(v)$	SFE4

– if $\varrho_m^\pi(s) = \mathsf{f}$ for each $\pi \in \mathcal{P}$, s is not able to process method m and proceeds as δ.

The empty service δ itself is unable to process any method. A service is fully deterministic if, for all m, for all s, $\varrho_m^\pi(s) = \mathsf{t}$ only if $\pi \in \{0, 1\}$.

The assumptions with respect to services made above are the ones made before for the non-probabilistic case in e.g. [15] adapted to the probabilistic case.

It is also assumed that a fixed but arbitrary set $\mathcal{F}$ of *foci* has been given. Foci play the role of names of services in a service family.

SFA has the sorts, constants and operators from $\Sigma_{\mathcal{S}}$ and in addition the sort **SF** of *service families* and the following constant and operators:

– the *empty service family* constant $\emptyset : \to \mathbf{SF}$;
– for each $f \in \mathcal{F}$, the unary *singleton service family* operator $f._ : \mathbf{S} \to \mathbf{SF}$;
– the binary *service family composition* operator $_ \oplus _ : \mathbf{SF} \times \mathbf{SF} \to \mathbf{SF}$;
– for each $F \subseteq \mathcal{F}$, the unary *encapsulation* operator $\partial_F : \mathbf{SF} \to \mathbf{SF}$.

We assume that there are infinitely many variables of sort **S**, including s, and infinitely many variables of sort **SF**, including u, v, w. Terms are built as usual in the many-sorted case (see e.g. [31,34]). We use prefix notation for the singleton service family operators and infix notation for the service family composition operator.

The service family denoted by $\emptyset$ is the empty service family. The service family denoted by a closed term of the form $f.t$ consists of one named service only, the service concerned is the service denoted by t, and the name of this service is f. The service family denoted by a closed term of the form $t \oplus t'$ consists of all named services that belong to either the service family denoted by t or the service family denoted by t'. In the case where a named service from the service family denoted by t and a named service from the service family denoted by t' have the same name, they collapse to an empty service with the name concerned. The service family denoted by a closed term of the form $\partial_F(t)$ consists of all named services with a name not in F that belong to the service family denoted by t.

The axioms of SFA are given in Table 5. In this table, f stands for an arbitrary focus from $\mathcal{F}$ and F stands for an arbitrary subset of $\mathcal{F}$. These axioms simply formalize the informal explanation given above.

The constants, operators, and axioms of SFA were presented for the first time in [14].

Table 6. Axioms for the use operator

$\mathsf{D} \mathbin{/} u = \mathsf{D}$		prU1
$\mathsf{S} \mathbin{/} u = \mathsf{S}$		prU2
$(\mathsf{tau} \circ x) \mathbin{/} u = \mathsf{tau} \circ (x \mathbin{/} u)$		prU3
$(x \trianglelefteq f.m \trianglerighteq y) \mathbin{/} \partial_{\{f\}}(u) = (x \mathbin{/} \partial_{\{f\}}(u)) \trianglelefteq f.m \trianglerighteq (y \mathbin{/} \partial_{\{f\}}(u))$		prU4
$(x \trianglelefteq f.m \trianglerighteq y) \mathbin{/} (f.t \oplus \partial_{\{f\}}(u)) = \mathsf{tau} \circ ((x +_\pi y) \mathbin{/} (f.\frac{\partial}{\partial m}t \oplus \partial_{\{f\}}(u)))$		
	if $\varrho^\pi_m(t) = \mathsf{t}$	prU5
$(x \trianglelefteq f.m \trianglerighteq y) \mathbin{/} (f.t \oplus \partial_{\{f\}}(u)) = \mathsf{tau} \circ \mathsf{D}$	if $\bigwedge_{\pi \in \mathcal{P}} \varrho^\pi_m(t) = \mathsf{f}$	prU6
$(x +_\pi y) \mathbin{/} u = (x \mathbin{/} u) +_\pi (y \mathbin{/} u)$		prU7

Table 7. Axioms for the abstraction operator

$\tau_{\mathsf{tau}}(\mathsf{S}) = \mathsf{S}$	TA1
$\tau_{\mathsf{tau}}(\mathsf{D}) = \mathsf{D}$	TA2
$\tau_{\mathsf{tau}}(\mathsf{tau} \circ x) = \tau_{\mathsf{tau}}(x)$	TA3
$\tau_{\mathsf{tau}}(x \trianglelefteq f.m \trianglerighteq y) = \tau_{\mathsf{tau}}(x) \trianglelefteq f.m \trianglerighteq \tau_{\mathsf{tau}}(y)$	TA4
$\tau_{\mathsf{tau}}(x +_\pi y) = \tau_{\mathsf{tau}}(x) +_\pi \tau_{\mathsf{tau}}(y)$	TA5

For the set $\mathcal{A}$ of basic actions, we now take $\{f.m \mid f \in \mathcal{F}, m \in \mathcal{M}\}$. Performing a basic action $f.m$ is taken as making a request to the service named f to process method m.

We combine prBTA with SFA and extend the combination with the following operators:

- the binary *use* operator $_ \mathbin{/} _ : \mathbf{T} \times \mathbf{SF} \to \mathbf{T}$;
- the unary *abstraction* operator $\tau_{\mathsf{tau}} : \mathbf{T} \to \mathbf{T}$;

and the axioms given in Tables 6 and 7, and call the resulting theory $\mathrm{prTA}_{\mathrm{tsi}}$. In these tables, f stands for an arbitrary focus from $\mathcal{F}$, m stands for an arbitrary method from $\mathcal{M}$, π stands for an arbitrary probability from $\mathcal{P}$, and t stands for an arbitrary term of sort $\mathbf{S}$. The axioms formalize the informal explanation given below. We use infix notation for the use operator.

The thread denoted by a closed term of the form $t \mathbin{/} t'$ is the thread that results from processing the method of each basic action performed by the thread denoted by t by the service with the focus of the basic action as its name in the service family denoted by t' each time that a service with the name in question really exists and as long as the method concerned can be processed. In the case that a service with the name in question does not really exist, the processing of a method is simply skipped (axiom prU4). When the method of a basic action performed by the thread can be processed by the named service, that service changes in accordance with the method and the thread is affected as follows: the basic action is turned into the internal action tau and then an internal choice is made between the two ways to proceed according to the probabilities of the two

possible reply values in the case of the method concerned (axiom prU5). When the method of a basic action performed by the thread cannot be processed by the named service, inaction occurs after the basic action is turned into the internal action tau (axiom prU6).

The thread denoted by a closed term of the form $\tau_{\mathsf{tau}}(t)$ is the thread that results from concealing the presence of the internal action tau in the thread denoted by t.

The use operator and the abstraction operator are adopted from the extension of BTA with thread-service interaction presented before in [15]. With the exception of axiom prU7, the axioms for the use operator are the ones given before for the non-probabilistic case in [15] adapted to the probabilistic case. With the exception of axiom TA5, the axioms for the abstraction operator are adopted from the extension of BTA with thread-service interaction presented in [15]. Axiom prU7 and TA5 are new.

The following theorem concerns the question whether the operators added to prBTA in $\mathrm{prTA}_{\mathrm{tsi}}$ are well axiomatized by the equations given in Tables 6 and 7 in the sense that these equations allow the projective limit model of prBTA to be expanded to a projective limit model of $\mathrm{prTA}_{\mathrm{tsi}}$.

Theorem 1. *The operators added to* prBTA *are well axiomatized, i.e.:*

(a) for all closed $\mathrm{prTA}_{\mathrm{tsi}}$ *terms* t *of sort* $\mathbf{T}$*, there exists a closed* prBTA *term* t' *such that* $t = t'$ *is derivable from the axioms of* $\mathrm{prTA}_{\mathrm{tsi}}$*;*

(b) for all closed prBTA *terms* t *and* t'*,* $t = t'$ *is derivable from the axioms of* prBTA *iff* $t = t'$ *is derivable from the axioms of* $\mathrm{prTA}_{\mathrm{tsi}}$*;*

(c) for all closed $\mathrm{prTA}_{\mathrm{tsi}}$ *terms* t *of sort* $\mathbf{T}$*, closed* $\mathrm{prTA}_{\mathrm{tsi}}$ *terms* t' *of sort* $\mathbf{SF}$ *and* $n \in \mathbb{N}$*,* $\pi_n(t \,/\, t') = \pi_n(\pi_n(t) \,/\, t')$ *is derivable from the axioms of* $\mathrm{prTA}_{\mathrm{tsi}}$ *and the following axioms for the unary operators* π_n *(which are explained below):*[2]

$$\begin{aligned}
&\pi_0(x) = \mathsf{D}\;, \\
&\pi_{n+1}(\mathsf{D}) = \mathsf{D}\;, \qquad \pi_{n+1}(x \trianglelefteq a \trianglerighteq y) = \pi_n(x) \trianglelefteq a \trianglerighteq \pi_n(y)\;, \\
&\pi_{n+1}(\mathsf{S}) = \mathsf{S}\;, \qquad \pi_{n+1}(x +_\pi y) = \pi_{n+1}(x) +_\pi \pi_{n+1}(y)\;.
\end{aligned}$$

where n *stands for an arbitrary natural number from* $\mathbb{N}$*,* a *stands for an arbitrary basic action from* $\mathcal{A}_{\mathsf{tau}}$*, and* π *is an arbitrary probability from* $\mathcal{P}$*;*

(d) for all closed $\mathrm{prTA}_{\mathrm{tsi}}$ *terms* t *of sort* $\mathbf{T}$ *and* $n \in \mathbb{N}$*, there exists a* $k \in \mathbb{N}$ *such that, for all* $m \in \mathbb{N}$ *with* $m \geq k$*,* $\pi_n(\tau_{\mathsf{tau}}(t)) = \pi_n(\tau_{\mathsf{tau}}(\pi_m(t)))$ *is derivable from the axioms of* $\mathrm{prTA}_{\mathrm{tsi}}$ *and the axioms for the operators* π_n *introduced in part (c).*

Proof. Part (a) is easily proved by induction on the structure of t, and in the case where t is of the form $t_1 \,/\, t_2$ and the case where t is of the form $\tau_{\mathsf{tau}}(t_1)$ by induction on the structure of t_1. In the subcase where t is of the form $t'_1 \trianglelefteq a \trianglerighteq t'_1 \,/\, t_2$,

[2] Holding on to the usual conventions leads to the double use of the symbol π: without subscript it stands for a probability value and with subscript it stands for a projection operator.

we need the easy to prove fact that, for each $f \in \mathcal{F}$ and closed term t of sort **SF**, either $t = \partial_f(t)$ is derivable or there exists a closed term t' of sort **S** such that $t = f.t' \oplus \partial_f(t)$ is derivable.

In the case of part (b), the implication from left to right follows immediately from the fact that the axioms of prBTA are included in the axioms of $\mathrm{prTA}_{\mathrm{tsi}}$. The implication from right to left is not difficult to see either. From the axioms of $\mathrm{prTA}_{\mathrm{tsi}}$ that are not axioms of prBTA, only axioms prU1, prU2, prU6, TA1, and TA2 may be applicable to a closed prBTA term t. If one of them is applicable, then the application yields an equation $t = t'$ in which t' is not a closed prBTA term. Moreover, only the axiom whose application yielded $t = t'$ is applicable to t', but now in the opposite direction. Hence, applications of axioms of $\mathrm{prTA}_{\mathrm{tsi}}$ that are not axioms of prBTA do not yield additional equations.

By part (a), it is sufficient to prove parts (c) and (d) for all closed prBTA terms t. Parts (c) and (d) are easily proved by induction on the structure of t, and in each case by case distinction between $n = 0$ and $n > 0$. In the proof of both parts, we repeatedly need the easy to prove fact that, for all closed prBTA terms t and $n \in \mathbb{N}$, $\pi_n(t) = \pi_n(\pi_n(t))$ is derivable. In the proof of part (c), in the case where t is of the form $t_1 \trianglelefteq a \trianglerighteq t_2$, we need again the fact mentioned at the end of the proof outline of part (a). □

The unary operators π_n are called *projection* operators. The thread denoted by a closed term of the form $\pi_n(t)$ is the thread that differs from the thread denoted by t in that it becomes inactive as soon as it has performed n actions.

By parts (a) and (b) of Theorem 1, we know that the carrier of the projective limit model of prBTA can serve as the carrier of a projective limit model of $\mathrm{prTA}_{\mathrm{tsi}}$ if it is possible to define on this carrier operations corresponding to the added operators such that the added equations are satisfied. By parts (c) and (d) of Theorem 1, we know that it is possible to do so. Thus, we know that the projective limit model of prBTA can be expanded to a projective limit model of $\mathrm{prTA}_{\mathrm{tsi}}$.

The actual expansion goes along the same lines as in the non-probabilistic case (see [15]). An outline of this expansion is given in Appendix A.2. Because the depth of the approximations of a thread may decrease by abstraction, we do not have that, for all n and t, $\pi_n(\tau_{\mathsf{tau}}(t)) = \pi_n(\tau_{\mathsf{tau}}(\pi_n(t)))$ is derivable. However, it is sufficient that there exists a $k \in \mathbb{N}$ such that, for all $m \in \mathbb{N}$ with $m \geq k$, $\pi_n(\tau_{\mathsf{tau}}(t)) = \pi_n(\tau_{\mathsf{tau}}(\pi_m(t)))$ is derivable (see also [15]).

5 A Probabilistic Program Notation

In this section, we introduce the probabilistic program notation prPGLB (probabilistic PGLB). In [6], a hierarchy of program notations rooted in program algebra is presented. One of the program notations that belong to this hierarchy is PGLB (ProGramming Language B). This program notation is close to existing assembly languages and has relative jump instructions. The program notation prPGLB is PGLB extended with probabilistic instructions that allow probabilistic choices to be made during the execution of instruction sequences.

In prPGLB, it is assumed that a fixed but arbitrary non-empty finite set $\mathfrak{A}$ of *basic instructions* has been given. The intuition is that the execution of a basic instruction in most instances modifies a state and in all instances produces a reply at its completion. The possible replies are the values t and f, and the actual reply is in most instances state-dependent. Therefore, successive executions of the same basic instruction may produce different replies. The set $\mathfrak{A}$ is the basis for the set of all instructions that may appear in the instruction sequences considered in prPGLB. These instructions are called primitive instructions.

The program notation prPGLB has the following primitive instructions:

- for each $a \in \mathfrak{A}$, a *plain basic instruction* a;
- for each $a \in \mathfrak{A}$, a *positive test instruction* $+a$;
- for each $a \in \mathfrak{A}$, a *negative test instruction* $-a$;
- for each $\pi \in \mathcal{P}$, a *plain random choice instruction* $\%(\pi)$;
- for each $\pi \in \mathcal{P}$, a *positive random choice instruction* $+\%(\pi)$;
- for each $\pi \in \mathcal{P}$, a *negative random choice instruction* $-\%(\pi)$;
- for each $l \in \mathbb{N}$, a *forward jump instruction* $\#l$;
- for each $l \in \mathbb{N}$, a *backward jump instruction* $\backslash\#l$;
- a *termination instruction* $!$.

A prPGLB instruction sequence has the form $u_1 \,;\ldots;\, u_k$, where $u_1, \ldots, u_k$ are primitive instructions of prPGLB.

On execution of a prPGLB instruction sequence, these primitive instructions have the following effects:

- the effect of a positive test instruction $+a$ is that basic instruction a is executed and execution proceeds with the next primitive instruction if t is produced and otherwise the next primitive instruction is skipped and execution proceeds with the primitive instruction following the skipped one – if there is no primitive instruction to proceed with, execution becomes inactive;
- the effect of a negative test instruction $-a$ is the same as the effect of $+a$, but with the role of the value produced reversed;
- the effect of a plain basic instruction a is the same as the effect of $+a$, but execution always proceeds as if t is produced;
- the effect of a positive random choice instruction $+\%(\pi)$ is that first t is produced with probability π and f is produced with probability $1-\pi$ and then execution proceeds with the next primitive instruction if t is produced and otherwise the next primitive instruction is skipped and execution proceeds with the primitive instruction following the skipped one – if there is no primitive instruction to proceed with, execution becomes inactive;
- the effect of a negative random choice instruction $-\%(\pi)$ is the same as the effect of $+\%(\pi)$, but with the role of the value produced reversed;
- the effect of a plain random choice instruction $\%(\pi)$ is the same as the effect of $+\%(\pi)$, but execution always proceeds as if t is produced;
- the effect of a forward jump instruction $\#l$ is that execution proceeds with the l^{th} next primitive instruction – if l equals 0 or there is no primitive instruction to proceed with, execution becomes inactive;

– the effect of a backward jump instruction $\backslash\#l$ is that execution proceeds with the l^{th} previous primitive instruction – if l equals 0 or there is no primitive instruction to proceed with, execution becomes inactive;
– the effect of the termination instruction ! is that execution terminates.

With the exception of the random choice instructions, the primitive instructions of prPGLB are adopted from PGLB. Counterparts of the random choice instructions are especially found in probabilistic extensions of Dijkstra's guarded command language (see e.g. [22]).

In order to describe the behaviours produced by prPGLB instruction sequences on execution, we need a service that behaves as a random Boolean generator. This service is able to process the following methods:

– for each $\pi \in \mathcal{P}$, a *get random Boolean method* $\mathsf{get}(\pi)$.

For each $\pi \in \mathcal{P}$, the method $\mathsf{get}(\pi)$ can be explained as follows: the service produces the reply t with probability π and the reply f with probability $1 - \pi$.

For the carrier of sort $\mathbf{S}$, we take the set $\{RBG, \delta\}$. For each $m \in \mathcal{M}$ and $\pi \in \mathcal{P}$, we take the functions $\frac{\partial}{\partial m}$ and ϱ^{π}_{m} such that:

$$\frac{\partial}{\partial \mathsf{get}(\pi)}(RBG) = RBG\ , \qquad \frac{\partial}{\partial m}(RBG) = \delta \text{ if } m \notin \{\mathsf{get}(\pi) \mid \pi \in \mathcal{P}\}\ ,$$
$$\varrho^{\pi}_{\mathsf{get}(\pi)}(RBG) = \mathsf{t}\ , \qquad \varrho^{\pi}_{m}(RBG) = \mathsf{f} \text{ if } m \neq \mathsf{get}(\pi)\ .$$

Moreover, we take the name RBG used above to denote the element of the carrier of sort $\mathbf{S}$ that differs from δ for a constant of sort $\mathbf{S}$. It is assumed that $\mathsf{get}(\pi) \in \mathcal{M}$ for each $\pi \in \mathcal{P}$. It is also assumed that $\mathsf{rbg} \in \mathcal{F}$.

The behaviours produced by prPGLB instruction sequences on execution are considered to be probabilistic threads, with the basic instructions taken as basic actions. The *thread extraction* operation $|_|$ defines, for each prPGLB instruction sequence, the behaviour produced on its execution. The thread extraction operation is defined by

$$|u_1\,;\ldots;\,u_k| = \tau_{\mathsf{tau}}(|1, u_1\,;\ldots;\,u_k| \,/\, \mathsf{rbg}.RBG)\ ,$$

where $|_,_|$ is defined by the equations given in Table 8 (for $a \in \mathfrak{A}$, $\pi \in \mathcal{P}$, and $l, i \in \mathbb{N}$)[3] and the rule that $|i, u_1\,;\ldots;\,u_k| = \mathsf{D}$ if u_i is the beginning of an infinite jump chain.[4]

If $1 \leq i \leq k$, $\tau_{\mathsf{tau}}(|i, u_1\,;\ldots;\,u_k| \,/\, \mathsf{rbg}.RBG)$ can be read as the behaviour produced by $u_1\,;\ldots;\,u_k$ on execution if execution starts at the i^{th} primitive instruction. By default, execution starts at the first primitive instruction.

In [12], we proposed several kinds of probabilistic jump instructions (bounded and unbounded, according to uniform probability distributions and geometric probability distributions). The meaning of instruction sequences from extensions of prPGLB with these kinds of probabilistic instructions can be given by a translation to instruction sequences from prPGLB.

[3] We write $i \dot{-} j$ for the monus of i and j, i.e. $i \dot{-} j = i - j$ if $i \geq j$ and $i \dot{-} j = 0$ otherwise.
[4] This rule can be formalized, cf. [11].

Table 8. Defining equations for the thread extraction operation

$\lvert i, u_1 ; \ldots ; u_k\rvert = \mathsf{D}$	if $\neg\, 1 \leq i \leq k$
$\lvert i, u_1 ; \ldots ; u_k\rvert = a \circ \lvert i+1, u_1 ; \ldots ; u_k\rvert$	if $u_i = a$
$\lvert i, u_1 ; \ldots ; u_k\rvert = \lvert i+1, u_1 ; \ldots ; u_k\rvert \trianglelefteq a \trianglerighteq \lvert i+2, u_1 ; \ldots ; u_k\rvert$	if $u_i = +a$
$\lvert i, u_1 ; \ldots ; u_k\rvert = \lvert i+2, u_1 ; \ldots ; u_k\rvert \trianglelefteq a \trianglerighteq \lvert i+1, u_1 ; \ldots ; u_k\rvert$	if $u_i = -a$
$\lvert i, u_1 ; \ldots ; u_k\rvert = \mathsf{rbg.get}(\pi) \circ \lvert i+1, u_1 ; \ldots ; u_k\rvert$	if $u_i = \%(\pi)$
$\lvert i, u_1 ; \ldots ; u_k\rvert = \lvert i+1, u_1 ; \ldots ; u_k\rvert \trianglelefteq \mathsf{rbg.get}(\pi) \trianglerighteq \lvert i+2, u_1 ; \ldots ; u_k\rvert$	if $u_i = +\%(\pi)$
$\lvert i, u_1 ; \ldots ; u_k\rvert = \lvert i+2, u_1 ; \ldots ; u_k\rvert \trianglelefteq \mathsf{rbg.get}(\pi) \trianglerighteq \lvert i+1, u_1 ; \ldots ; u_k\rvert$	if $u_i = -\%(\pi)$
$\lvert i, u_1 ; \ldots ; u_k\rvert = \lvert i+l, u_1 ; \ldots ; u_k\rvert$	if $u_i = \#l$
$\lvert i, u_1 ; \ldots ; u_k\rvert = \lvert i \dot{-} l, u_1 ; \ldots ; u_k\rvert$	if $u_i = \backslash\#l$
$\lvert i, u_1 ; \ldots ; u_k\rvert = \mathsf{S}$	if $u_i = !$

6 Probabilistic Strategic Interleaving of Threads

Multi-threading refers to the concurrent existence of several threads in a program under execution. It is the dominant form of concurrency provided by contemporary programming languages such as Java [21] and C# [23]. Theories of concurrent processes such as ACP [3], CCS [27], and CSP [24] are based on arbitrary interleaving. In the case of multi-threading, more often than not some interleaving strategy is used. We abandon the point of view that arbitrary interleaving is the most appropriate abstraction when dealing with multi-threading. The following points illustrate why we find difficulty in taking that point of view: (a) whether the interleaving of certain threads leads to inactiveness depends on the interleaving strategy used; (b) sometimes inactiveness occurs with a particular interleaving strategy whereas arbitrary interleaving would not lead to inactiveness, and vice versa. Demonstrations of (a) and (b) are given in [8] and [7], respectively.

The probabilistic features of prBTA allow it to be extended with interleaving strategies that correspond to probabilistic scheduling algorithms. In this section, we take up the extension of prBTA with such probabilistic interleaving strategies. The presented extension covers an arbitrary probabilistic interleaving strategy that can be represented in the way that is explained below.

We write $\mathcal{A}'_{\mathsf{tau}}$ for $\mathcal{A}_{\mathsf{tau}} \cup \{\mathsf{nt}, \mathsf{S}, \mathsf{D}\}$ and we write $\mathcal{H}$ for $(\mathbb{N}_1 \times \mathbb{N}_1)^*$.[5] The elements of $\mathcal{H}$ are called *interleaving histories*. The intuition concerning interleaving histories is as follows: if the jth pair of an interleaving history is (i, n), then the ith thread got a turn in the jth interleaving step and after its turn there were n threads to be interleaved.

With regard to interleaving of threads, it is assumed that the following has been given:

- a set S;

[5] We write $\mathbb{N}_1$ for the set $\{n \in \mathbb{N} \mid n \geq 1\}$ of positive natural numbers.

– an indexed family of functions $\langle \sigma_n \rangle_{n \in \mathbb{N}_1}$ where $\sigma_n : \mathcal{H} \times S \to (\{1, \ldots, n\} \to \mathcal{P})$ for each $n \in \mathbb{N}_1$;
– an indexed family of functions $\langle \vartheta_n \rangle_{n \in \mathbb{N}_1}$ where $\vartheta_n : \mathcal{H} \times S \times \{1, \ldots, n\} \times \mathcal{A}'_{\mathsf{tau}} \to S$ for each $n \in \mathbb{N}_1$.

The elements of S are called *control states*, σ_n is called an *abstract scheduler* (for n threads), and ϑ_n is called a *control state transformer* (for n threads). The intuition concerning S, $\langle \sigma_n \rangle_{n \in \mathbb{N}_1}$, and $\langle \vartheta_n \rangle_{n \in \mathbb{N}_1}$ is as follows:

– the control states from S encode data relevant to the interleaving strategy (e.g., for each of the threads being interleaved, the set of all foci naming services on which it currently keeps a lock);
– for each $h \in \mathcal{H}$ and $s \in S$, $\sigma_n(h, s)$ is the probability distribution on n threads that assigns to each of the threads the probability that it gets the next turn after history h in state s;
– for each $h \in \mathcal{H}$, $s \in S$, $i \in \{1, \ldots, n\}$, and $a \in \mathcal{A}'_{\mathsf{tau}}$, $\vartheta_n(h, s, i, a)$ is the control state that arises after history h in state s on the ith thread doing a.

Thus, S, $\langle \sigma_n \rangle_{n \in \mathbb{N}_1}$, and $\langle \vartheta_n \rangle_{n \in \mathbb{N}_1}$ provide a way to represent a probabilistic interleaving strategy. The abstraction of a scheduler used here is essentially the notion of a scheduler defined in [30].

We extend prBTA with the following operators:

– the ternary *forking postconditional composition* operator $_ \trianglelefteq \mathsf{nt}(_) \trianglerighteq _ : \mathbf{T} \times \mathbf{T} \times \mathbf{T} \to \mathbf{T}$;
– for each $n \in \mathbb{N}_1$, $h \in \mathcal{H}$, and $s \in S$, the n-ary *strategic interleaving* operator $\|^n_{h,s} : \mathbf{T} \times \cdots \times \mathbf{T} \to \mathbf{T}$;
– for each $n, i \in \mathbb{N}_1$ with $i \leq n$, $h \in \mathcal{H}$, and $s \in S$, the n-ary *positional strategic interleaving* operator $\between^{n,i}_{h,s} : \mathbf{T} \times \cdots \times \mathbf{T} \to \mathbf{T}$;
– the unary *deadlock at termination* operator $\mathsf{S_D} : \mathbf{T} \to \mathbf{T}$;

and the axioms given in Table 9,[6] and call the resulting theory $\mathrm{prTA}_{\mathrm{si}}$. In this table, n and i stand for arbitrary numbers from $\mathbb{N}_1$ with $i \leq n$, h stands for an arbitrary interleaving history from $\mathcal{H}$, s stands for an arbitrary control state from S, a stands for an arbitrary basic action from $\mathcal{A}_{\mathsf{tau}}$, and π stands for an arbitrary probability from $\mathcal{P}$.

The forking postconditional composition operator has the same shape as the postconditional composition operators introduced in Section 3. Formally, no basic action is involved in forking postconditional composition. However, for an operational intuition, in $t \trianglelefteq \mathsf{nt}(t'') \trianglerighteq t'$, $\mathsf{nt}(t'')$ can be considered a thread forking action. It represents the act of forking off thread t''. Like with real basic actions, a reply is produced upon performing a thread forking action.

The thread denoted by a closed term of the form $\|^n_{h,s}(t_1, \ldots, t_n)$ is the thread that results from interleaving of the n threads denoted by $t_1, \ldots, t_n$ after history

[6] We write $\langle\,\rangle$ for the empty sequence, d for the sequence having d as sole element, and $\alpha \frown \alpha'$ for the concatenation of sequences α and α'. We assume that the usual identities, such as $\langle\,\rangle \frown \alpha = \alpha$ and $(\alpha \frown \alpha') \frown \alpha'' = \alpha \frown (\alpha' \frown \alpha'')$, hold.

Table 9. Axioms for strategic interleaving

$\lVert^n_{h,s}(x_1,\ldots,x_n) = \sum_{i=1}^{n}[\sigma_n(h,s)(i)]\ \between^{n,i}_{h,s}(x_1,\ldots,x_n)$	prSI1
$\between^{1,i}_{h,s}(\mathsf{D}) = \mathsf{D}$	prSI2
$\between^{n+1,i}_{h,s}(x_1,\ldots,x_{i-1},\mathsf{D},x_{i+1},\ldots,x_{n+1}) =$ $\mathsf{S}_\mathsf{D}(\lVert^n_{h\frown(i,n),\vartheta_{n+1}(h,s,i,\mathsf{D})}(x_1,\ldots,x_{i-1},x_{i+1},\ldots,x_{n+1}))$	prSI3
$\between^{1,i}_{h,s}(\mathsf{S}) = \mathsf{S}$	prSI4
$\between^{n+1,i}_{h,s}(x_1,\ldots,x_{i-1},\mathsf{S},x_{i+1},\ldots,x_{n+1}) =$ $\lVert^n_{h\frown(i,n),\vartheta_{n+1}(h,s,i,\mathsf{S})}(x_1,\ldots,x_{i-1},x_{i+1},\ldots,x_{n+1})$	prSI5
$\between^{n,i}_{h,s}(x_1,\ldots,x_{i-1},x'_i \trianglelefteq \mathsf{nt}(x) \trianglerighteq x''_i,x_{i+1},\ldots,x_n) =$ $\mathsf{tau}\circ\lVert^{n+1}_{h\frown(i,n+1),\vartheta_n(h,s,i,\mathsf{nt})}(x_1,\ldots,x_{i-1},x'_i,x_{i+1},\ldots,x_n,x)$	prSI6
$\between^{n,i}_{h,s}(x_1,\ldots,x_{i-1},x'_i \trianglelefteq a \trianglerighteq x''_i,x_{i+1},\ldots,x_n) =$ $\lVert^n_{h\frown(i,n),\vartheta_n(h,s,i,a)}(x_1,\ldots,x_{i-1},x'_i,x_{i+1},\ldots,x_n)$ $\trianglelefteq a \trianglerighteq$ $\lVert^n_{h\frown(i,n),\vartheta_n(h,s,i,a)}(x_1,\ldots,x_{i-1},x''_i,x_{i+1},\ldots,x_n)$	prSI7
$\between^{n,i}_{h,s}(x_1,\ldots,x_{i-1},x'_i +_\pi x''_i,x_{i+1},\ldots,x_n) =$ $\between^{n,i}_{h,s}(x_1,\ldots,x_{i-1},x'_i,x_{i+1},\ldots,x_n)$ $+_\pi$ $\between^{n,i}_{h,s}(x_1,\ldots,x_{i-1},x''_i,x_{i+1},\ldots,x_n)$	prSI8
$\mathsf{S}_\mathsf{D}(\mathsf{D}) = \mathsf{D}$	DT1
$\mathsf{S}_\mathsf{D}(\mathsf{S}) = \mathsf{D}$	DT2
$\mathsf{S}_\mathsf{D}(x \trianglelefteq \mathsf{nt}(z) \trianglerighteq y) = \mathsf{S}_\mathsf{D}(x) \trianglelefteq \mathsf{nt}(\mathsf{S}_\mathsf{D}(z)) \trianglerighteq \mathsf{S}_\mathsf{D}(y)$	DT3
$\mathsf{S}_\mathsf{D}(x \trianglelefteq a \trianglerighteq y) = \mathsf{S}_\mathsf{D}(x) \trianglelefteq a \trianglerighteq \mathsf{S}_\mathsf{D}(y)$	DT4
$\mathsf{S}_\mathsf{D}(x +_\pi y) = \mathsf{S}_\mathsf{D}(x) +_\pi \mathsf{S}_\mathsf{D}(y)$	DT5

h in state s, according to the interleaving strategy represented by S, $\langle\sigma_n\rangle_{n\in\mathbb{N}_1}$, and $\langle\vartheta_n\rangle_{n\in\mathbb{N}_1}$. By the interleaving, a number of threads is turned into a single thread. In this single thread, the internal action tau arises as a residue of each thread forking action encountered. Moreover, the possibility that f is produced as a reply upon performing a thread forking action is ignored. This reflects our focus on the case where capacity problems with respect to thread forking never arise.

The positional strategic interleaving operators are auxiliary operators used to axiomatize the strategic interleaving operators. The role of the positional strategic interleaving operators in the axiomatization is similar to the role of the left merge operator found in process algebra (see e.g. [3]). The deadlock at termination operator is an auxiliary operator as well. It is used in axiom prSI3 to express that in the event of inactiveness of one thread, the whole become inactive only after all other threads have terminated or become inactive. The thread denoted by a closed term of the form $\mathsf{S}_\mathsf{D}(t)$ is the thread that results from turning termination into inactiveness in the thread denoted by t.

The forking postconditional composition operator and the deadlock at termination operator are adopted from earlier extensions of BTA with strategic interleaving. The strategic interleaving operators and the positional strategic interleaving operators are not adopted from earlier extensions of BTA with strategic interleaving. To our knowledge, no probabilistic process algebras with counterparts of these operators has been proposed until now. Axioms prSI1–prSI8 and DT5 are new. Axioms DT1–DT4 are adopted from the extension of BTA with strategic interleaving and thread forking presented in [13].

Consider the case where S is a singleton set, for each $n \in \mathbb{N}_1$, σ_n is defined by

$$\begin{aligned}
&\sigma_n(\langle\,\rangle, s)(i) = 1 && \text{if } i = 1\ ,\\
&\sigma_n(\langle\,\rangle, s)(i) = 0 && \text{if } i \neq 1\ ,\\
&\sigma_n(h \frown (j, n)\,, s)(i) = 1 && \text{if } i = (j+1) \bmod n\ ,\\
&\sigma_n(h \frown (i, n)\,, s)(i) = 0 && \text{if } i \neq (j+1) \bmod n
\end{aligned}$$

and, ϑ_n is defined by

$$\vartheta_n(h, s, i, a) = s\ .$$

In this case, the interleaving strategy corresponds to the round-robin scheduling algorithm. This deterministic interleaving strategy is called cyclic interleaving in our earlier work on interleaving strategies (see e.g. [8]). In the current setting, an interleaving strategy is deterministic if, for all n, for all h, s, and i, $\sigma_n(h, s)(i) \in \{0, 1\}$. In the case that S and ϑ_n are as above, but σ_n is defined by

$$\begin{aligned}
&\sigma_n(\langle\,\rangle, s)(i) = 1 && \text{if } i = 1\ ,\\
&\sigma_n(\langle\,\rangle, s)(i) = 0 && \text{if } i \neq 1\ ,\\
&\sigma_n(h \frown (j, n)\,, s)(i) = 1/n && \text{if } i \leq n\ ,\\
&\sigma_n(h \frown (i, n)\,, s)(i) = 0 && \text{if } i > n\ ,
\end{aligned}$$

the interleaving strategy is a purely probabilistic one. The probability distribution used is a uniform distribution.

More advanced strategies can be obtained if the scheduling makes use of the whole interleaving history and/or the control state. For example, the individual lifetimes of the threads to be interleaved and their creation hierarchy can be taken into account by making use of the whole interleaving history. Individual properties of the threads to be interleaved that depend on the actions performed by them can be taken into account by making use of the control state. By doing so, interleaving strategies are obtained which, to a certain extent, can be affected by the threads to be interleaved.

Henceforth, we will write $\mathrm{prBTA}_{\mathrm{nt}}$ for prBTA extended with the forking postconditional composition operator. The projective limit model of $\mathrm{prBTA}_{\mathrm{nt}}$ is constructed like the projective limit model of prBTA. An outline of the projective limit model of $\mathrm{prBTA}_{\mathrm{nt}}$ is given in Appendix A.3.

The following theorem concerns the question whether the operators added to $\mathrm{prBTA}_{\mathrm{nt}}$ are well axiomatized by the equations given in Table 9 in the sense

that these equations allow the projective limit model of $\mathrm{prBTA}_{\mathrm{nt}}$ to be expanded to a projective limit model of $\mathrm{prTA}_{\mathrm{si}}$.

Theorem 2. *The operators added to* $\mathrm{prBTA}_{\mathrm{nt}}$ *are well axiomatized, i.e.:*

(a) for all closed $\mathrm{prTA}_{\mathrm{si}}$ *terms* t*, there exists a closed* $\mathrm{prBTA}_{\mathrm{nt}}$ *term* t' *such that* $t = t'$ *is derivable from the axioms of* $\mathrm{prTA}_{\mathrm{si}}$*;*

(b) for all closed $\mathrm{prBTA}_{\mathrm{nt}}$ *terms* t *and* t'*,* $t = t'$ *is derivable from the axioms of* $\mathrm{prBTA}_{\mathrm{nt}}$ *iff* $t = t'$ *is derivable from the axioms of* $\mathrm{prTA}_{\mathrm{si}}$*;*

(c) for all $m, i \in \mathbb{N}_1$ *with* $i \leq m$*,* $h \in \mathcal{H}$*,* $s \in S$*, closed* $\mathrm{prTA}_{\mathrm{si}}$ *terms* $t_1, \ldots, t_m$ *and* $n \in \mathbb{N}$*,* $\pi_n(\|^m_{h,s}(t_1, \ldots, t_m)) = \pi_n(\|^m_{h,s}(\pi_n(t_1), \ldots, \pi_n(t_m)))$ *and* $\pi_n(\between^{m,i}_{h,s}(t_1, \ldots, t_m)) = \pi_n(\between^{m,i}_{h,s}(\pi_n(t_1), \ldots, \pi_n(t_m)))$ *are derivable from the axioms of* $\mathrm{prTA}_{\mathrm{si}}$*, the axioms for the operators* π_n *introduced in Theorem 1, and the following axiom:*

$$\pi_{n+1}(x \trianglelefteq \mathsf{nt}(z) \trianglerighteq y) = \pi_{n+1}(x) \trianglelefteq \mathsf{nt}(\pi_{n+1}(z)) \trianglerighteq \pi_{n+1}(y) \;,$$

where n *stands for an arbitrary natural number from* $\mathbb{N}$*;*

(d) for all closed $\mathrm{prTA}_{\mathrm{si}}$ *terms* t *and* $n \in \mathbb{N}$*,* $\pi_n(\mathsf{S}_{\mathsf{D}}(t)) = \pi_n(\mathsf{S}_{\mathsf{D}}(\pi_n(t)))$ *is derivable from the axioms of* $\mathrm{prTA}_{\mathrm{si}}$*, the axioms for the operators* π_n *introduced in Theorem 1, and the axiom introduced in part (c).*

Proof. Part (a) is straightforwardly proved by induction on the structure of t, and then in the case where t is of the form $\between^{n,i}_{h,s}(t_1, \ldots, t_n)$ by induction on the sum of the lengths of $t_1, \ldots, t_n$ and case distinction on the structure of t_i and in the case where t is of the form $\mathsf{S}_{\mathsf{D}}(t_1)$ by induction on the structure of t_1. The proof of the case where t is of the form $\between^{n,i}_{h,s}(t_1, \ldots, t_n)$ reveals that occurrences of the forking postconditional composition operator get eliminated if t is of that form.

In the case of part (b), the implication from left to right follows immediately from the fact that the axioms of $\mathrm{prBTA}_{\mathrm{nt}}$ are included in the axioms of $\mathrm{prTA}_{\mathrm{si}}$. The implication from right to left is not difficult to see either. From the axioms of $\mathrm{prTA}_{\mathrm{si}}$ that are not axioms of $\mathrm{prBTA}_{\mathrm{nt}}$, only axioms prSI2, prSI4, DT1, and DT2 may be applicable to a closed $\mathrm{prBTA}_{\mathrm{nt}}$ term t. If one of them is applicable, then the application yields an equation $t = t'$ in which t' is not a closed $\mathrm{prBTA}_{\mathrm{nt}}$ term. Moreover, only the axiom whose application yielded $t = t'$ is applicable to t', but now in the opposite direction. Hence, applications of axioms of $\mathrm{prTA}_{\mathrm{si}}$ that are not axioms of $\mathrm{prBTA}_{\mathrm{nt}}$ do not yield additional equations.

By part (a), it is sufficient to prove part (c) for all closed $\mathrm{prBTA}_{\mathrm{nt}}$ terms $t_1, \ldots, t_m$. The derivability of the second equation is straightforwardly proved by induction on the sum of the lengths of $t_1, \ldots, t_n$ and case distinction on the structure of t_i, and in each case by case distinction between $n = 0$ and $n > 0$. The derivability of the first equation now follows immediately using the axioms of the operators π_n. In the proofs, we repeatedly need the easy to prove fact that, for all closed $\mathrm{prBTA}_{\mathrm{nt}}$ terms t and $n \in \mathbb{N}$, $\pi_n(t) = \pi_n(\pi_n(t))$ is derivable.

By part (a), it is sufficient to prove part (d) for all closed $\mathrm{prBTA}_{\mathrm{nt}}$ terms t. Part (d) is easily proved by induction on the structure of t, and in each case by

case distinction between $n = 0$ and $n > 0$. In the proof, we need again the fact mentioned at the end of the proof outline of part (c). □

By Theorem 2, we know that the projective limit model of $\mathrm{prBTA}_{\mathrm{nt}}$ can be expanded to a projective limit model of $\mathrm{prTA}_{\mathrm{si}}$. An outline of this expansion is given in Appendix A.3.

7 Concluding Remarks

We have added probabilistic features to BTA and its extensions with thread-service interaction and strategic interleaving. Thus, we have paved the way for rigorous investigation of issues related to probabilistic computation thinking in terms of instruction sequences and rigorous investigation of probabilistic interleaving strategies. As an example of the use of $\mathrm{prTA}_{\mathrm{tsi}}$, the probabilistic version of the extension of BTA with thread-service interaction, we have added the most basic kind of probabilistic instructions proposed in [12] to a program notation rooted in PGA and have given a formal definition of the behaviours produced by the instruction sequences from the resulting program notation under excution with the help of $\mathrm{prTA}_{\mathrm{tsi}}$.

We enumerate neither the numerous issues relating to probabilistic computation in areas such as computability and complexity of computational problems, efficiency of algorithms, and verification of programs that could be investigated thinking in terms of instruction sequences nor the numerous probabilistic scheduling algorithms that could be investigated in $\mathrm{prTA}_{\mathrm{si}}$, the probabilistic generalization of the extensions of BTA with strategic interleaving.

However, we mention interesting options for future work that are of a different kind: (a) clarifying analyses of relevant probabilistic algorithms, such as the Miller-Rabin probabilistic primality test [28], using probabilistic instruction sequences or non-probabilistic instruction sequences and probabilistic services and (b) explanations of relevant quantum algorithms, such as Shor's integer factorization algorithm [32], by first giving a clarifying analysis using probabilistic instruction sequences or non-probabilistic instruction sequences and probabilistic services and then showing how certain services involved in principle can be realized very efficiently with quantum computing.

Moreover, we believe that the development of program notations for probabilistic computation is a useful preparation for the development of program notations for quantum computation later on. The development of program notations for quantum computation that have their origins in instruction sequences could constitute a valuable complement to other developments with respect to quantum computation, which for the greater part boil down to mere adaptation of earlier developments with respect to classical computation to the potentialities of quantum physics (see e.g. [19]).

In fact, prBTA is a process algebra tailored to the behaviours produced by probabilistic instruction sequences under execution. Because prBTA offers probabilistic choices of the generative variety (see [20]) and no non-deterministic

choices, it is most closely related to the probabilistic process algebra prBPA presented in [2]. To our knowledge, thread-service interaction and strategic interleaving as found in $\mathrm{prTA}_{\mathrm{tsi}}$ and $\mathrm{prTA}_{\mathrm{si}}$ are mechanisms for interaction and concurrency that are quite different from those found in any theory or model of processes. This leaves almost nothing to be said about related work.

The very limited extent of related work is due to two conscious choices: (a) the limitation of the scope to behaviours produced by programs under execution and (b) the limitation of the scope to the form of interleaving concurrency that is relevant to the behaviours of multi-threaded programs under execution. However, something unexpected remains to be mentioned as related work, to wit the work on security of multi-threaded programs presented in [30]. Probabilistic strategic interleaving as found in $\mathrm{prTA}_{\mathrm{si}}$ is strongly inspired by the scheduler-dependent semantics of a simple programming language with support for multi-threading that we found in that paper.

It is noteworthy to mention something about the interpretation of prBTA, $\mathrm{prTA}_{\mathrm{tsi}}$, and $\mathrm{prTA}_{\mathrm{si}}$ in a probabilistic version of a general process algebra such as ACP, CCS or CSP. It is crucial that probabilistic choice of the generative variety, non-deterministic choice, asynchronous parallel composition, abstraction from internal actions, and recursion are covered by the process algebra used for the purpose of interpretation. General process algebras that cover all this are rare. To our knowledge, pACP_τ [1] is the only one that has been elaborated in sufficient depth. However, interpretation of prBTA, $\mathrm{prTA}_{\mathrm{tsi}}$, and $\mathrm{prTA}_{\mathrm{si}}$ in pACP_τ seems impossible to us. The presence of asynchronous parallel composition based on arbitrary interleaving in pACP_τ precludes the proper form of abstraction from internal actions for interpretation of prBTA, $\mathrm{prTA}_{\mathrm{tsi}}$, and $\mathrm{prTA}_{\mathrm{si}}$.

Acknowledgements

We thank two anonymous referees for carefully reading a preliminary version of this paper and for suggesting improvements of the presentation of the paper.

A Projective Limit Models

In this appendix, we outline the construction of projective limit models for prBTA, $\mathrm{prTA}_{\mathrm{tsi}}$, and $\mathrm{prTA}_{\mathrm{si}}$. In these model, which covers finite and infinite threads, threads are represented by infinite sequences of finite approximations. Guarded recursive specifications have unique solutions in these models. We denote the interpretations of constants and operators in the models by the constants and operators themselves.

A.1 Projective Limit Model of prBTA

We will write $\boldsymbol{\mathcal{I}}(\mathrm{prBTA})$ for the initial model of prBTA and $\boldsymbol{T}(\mathrm{prBTA})$ for the carrier of $\boldsymbol{\mathcal{I}}(\mathrm{prBTA})$. $\boldsymbol{T}(\mathrm{prBTA})$ consists of the equivalence classes of closed

prBTA terms with respect to derivable equality. In other words, modulo derivable equality, $\boldsymbol{T}(\mathrm{prBTA})$ is the set of all closed prBTA terms. Henceforth, we will identify closed prBTA terms with their equivalence class where elements of $\boldsymbol{T}(\mathrm{prBTA})$ are concerned.

Each element of $\boldsymbol{T}(\mathrm{prBTA})$ represents a finite thread, i.e. a thread with a finite upper bound to the number of actions that it can perform. Below, we will construct a model that covers infinite threads as well. In preparation for that, we define for all n a function that cuts off threads from $\boldsymbol{T}(\mathrm{prBTA})$ after n actions have been performed.

For each $n \in \mathbb{N}$, we define the *projection* function $\pi_n : \boldsymbol{T}(\mathrm{prBTA}) \to \boldsymbol{T}(\mathrm{prBTA})$, inductively as follows:

$$\begin{aligned}
&\pi_0(t) = \mathsf{D} \;, \\
&\pi_{n+1}(\mathsf{S}) = \mathsf{S} \;, \qquad \pi_{n+1}(t \trianglelefteq a \trianglerighteq t') = \pi_n(t) \trianglelefteq a \trianglerighteq \pi_n(t') \;, \\
&\pi_{n+1}(\mathsf{D}) = \mathsf{D} \;, \qquad \pi_{n+1}(t +_\pi t') = \pi_{n+1}(t) +_\pi \pi_{n+1}(t') \;.
\end{aligned}$$

For $t \in \boldsymbol{T}(\mathrm{prBTA})$, $\pi_n(t)$ is called the nth projection of t. It can be thought of as an approximation of t. If $\pi_n(t) \neq t$, then $\pi_{n+1}(t)$ can be thought of as the closest better approximation of t. If $\pi_n(t) = t$, then $\pi_{n+1}(t) = t$ as well. For all $n \in \mathbb{N}$, we will write $\boldsymbol{T}^n(\mathrm{prBTA})$ for $\{\pi_n(t) \mid t \in \boldsymbol{T}(\mathrm{prBTA})\}$. Obviously, the projection functions defined above satisfy the axioms for the projection operators introduced in Theorem 1.

In the projective limit model, which covers both finite and infinite threads, threads are represented by *projective sequences*, i.e. infinite sequences $(t_n)_{n\in\mathbb{N}}$ of elements of $\boldsymbol{T}(\mathrm{prBTA})$ such that $t_n \in \boldsymbol{T}^n(\mathrm{prBTA})$ and $t_n = \pi_n(t_{n+1})$ for all $n \in \mathbb{N}$. In other words, a projective sequence is a sequence of which successive components are successive projections of the same thread. The idea is that any infinite thread is fully characterized by the infinite sequence of all its finite approximations. We will write $\boldsymbol{T}^\infty(\mathrm{prBTA})$ for the set of all projective sequences over $\boldsymbol{T}(\mathrm{prBTA})$, i.e. the set

$$\{(t_n)_{n\in\mathbb{N}} \mid \textstyle\bigwedge_{n\in\mathbb{N}} (t_n \in \boldsymbol{T}^n(\mathrm{prBTA}) \wedge t_n = \pi_n(t_{n+1}))\} \;.$$

The *projective limit model* $\boldsymbol{\mathcal{I}}^\infty(\mathrm{prBTA})$ of prBTA consists of the following:

- the set $\boldsymbol{T}^\infty(\mathrm{prBTA})$, the carrier of the projective limit model;
- an element of $\boldsymbol{T}^\infty(\mathrm{prBTA})$ for each constant of prBTA;
- an operation on $\boldsymbol{T}^\infty(\mathrm{prBTA})$ for each operator of prBTA;

where those elements of $\boldsymbol{T}^\infty(\mathrm{prBTA})$ and operations on $\boldsymbol{T}^\infty(\mathrm{prBTA})$ are defined as follows:

$$\begin{aligned}
\mathsf{S} &= (\pi_n(\mathsf{S}))_{n\in\mathbb{N}} \;, \\
\mathsf{D} &= (\pi_n(\mathsf{D}))_{n\in\mathbb{N}} \;, \\
(t_n)_{n\in\mathbb{N}} \trianglelefteq a \trianglerighteq (t'_n)_{n\in\mathbb{N}} &= (\pi_n(t_n \trianglelefteq a \trianglerighteq t'_n))_{n\in\mathbb{N}} \;, \\
(t_n)_{n\in\mathbb{N}} +_\pi (t'_n)_{n\in\mathbb{N}} &= (\pi_n(t_n +_\pi t'_n))_{n\in\mathbb{N}} \;.
\end{aligned}$$

It is straightforward to check that the constants are elements of $\boldsymbol{T}^\infty(\text{prBTA})$ and the operations always yield elements of $\boldsymbol{T}^\infty(\text{prBTA})$. It follows immediately from the construction of the projective limit model of prBTA that the axiom of prBTA forms a complete axiomatization of this model for equations between closed terms.

A.2 Projective Limit Model of prTA_{tsi}

We will write $\boldsymbol{\mathcal{I}}(\text{SFA})$ for the free SFA-extension of $\boldsymbol{\mathcal{S}}$ and $\boldsymbol{\mathcal{I}}(\text{prTA}_{\text{tsi}})$ for the free prTA_{tsi}-extension of $\boldsymbol{\mathcal{S}}$.

From the fact that the signatures of $\boldsymbol{\mathcal{I}}^\infty(\text{prBTA})$ and $\boldsymbol{\mathcal{I}}(\text{SFA})$ are disjoint, it follows, by the amalgamation result about expansions presented as Theorem 6.1.1 in [25] (adapted to the many-sorted case), that there exists a model of prBTA combined with SFA such that the restriction to the signature of prBTA is $\boldsymbol{\mathcal{I}}^\infty(\text{prBTA})$ and the restriction to the signature of SFA is $\boldsymbol{\mathcal{I}}(\text{SFA})$.

Let $\boldsymbol{\mathcal{I}}^\infty(\text{prBTA+SFA})$ be the model of prBTA combined with SFA referred to above. Then the *projective limit model* $\boldsymbol{\mathcal{I}}^\infty(\text{prTA}_{\text{tsi}})$ of prTA_{tsi} is $\boldsymbol{\mathcal{I}}^\infty(\text{prBTA+SFA})$ expanded with the operations defined by

$$\begin{aligned}(t_n)_{n\in\mathbb{N}} \mathbin{/} S &= (\pi_n(t_n \mathbin{/} S))_{n\in\mathbb{N}}\,,\\ \tau_{\mathsf{tau}}((t_n)_{n\in\mathbb{N}}) &= (\lim_{k\to\infty} \pi_n(\tau_{\mathsf{tau}}(t_k)))_{n\in\mathbb{N}}\end{aligned}$$

as interpretations of the additional operators of prTA_{tsi}. On the right-hand side of these equations, the symbols / and τ_{tau} denote the interpretation of the operators / and τ_{tau} in $\boldsymbol{\mathcal{I}}(\text{prTA}_{\text{tsi}})$. In the second equation, the limit is the limit with respect to the discrete topology on $\boldsymbol{T}(\text{prBTA})$.

It is straightforward to check that the operations with which $\boldsymbol{\mathcal{I}}^\infty(\text{prBTA})$ is expanded always yield elements of $\boldsymbol{T}^\infty(\text{prBTA})$. It follows immediately from the construction of $\boldsymbol{\mathcal{I}}^\infty(\text{prTA}_{\text{tsi}})$ and Theorem 1 that $\boldsymbol{\mathcal{I}}^\infty(\text{prTA}_{\text{tsi}})$ is really a projective limit model of prTA_{tsi}.

A.3 Projective Limit Model of prTA_{si}

We will write $\boldsymbol{\mathcal{I}}(\text{prBTA}_{\text{nt}})$ for the initial model of prBTA_{nt} and $\boldsymbol{T}(\text{prBTA}_{\text{nt}})$ for the carrier of $\boldsymbol{\mathcal{I}}(\text{prBTA}_{\text{nt}})$. Moreover, we will write $\boldsymbol{\mathcal{I}}(\text{prTA}_{\text{si}})$ for the initial model of prTA_{si}.

With the projection functions π_n extended from $\boldsymbol{T}(\text{prBTA})$ to $\boldsymbol{T}(\text{prBTA}_{\text{nt}})$ such that

$$\pi_{n+1}(t \trianglelefteq \mathsf{nt}(t'') \trianglerighteq t') = \pi_{n+1}(t) \trianglelefteq \mathsf{nt}(\pi_{n+1}(t'')) \trianglerighteq \pi_{n+1}(t')\,,$$

the projective limit model $\boldsymbol{\mathcal{I}}^\infty(\text{prBTA}_{\text{nt}})$ of prBTA_{nt} is constructed from $\boldsymbol{\mathcal{I}}(\text{prBTA}_{\text{nt}})$ like the projective limit model $\boldsymbol{\mathcal{I}}^\infty(\text{prBTA})$ of prBTA is constructed from $\boldsymbol{\mathcal{I}}(\text{prBTA})$. The interpretation of the additional operator is the operation on $\boldsymbol{T}^\infty(\text{prBTA}_{\text{nt}})$ defined as follows:

$$(t_{1n})_{n\in\mathbb{N}} \trianglelefteq \mathsf{nt}((t_{2n})_{n\in\mathbb{N}}) \trianglerighteq (t_{3n})_{n\in\mathbb{N}} = (\pi_n(t_{1n} \trianglelefteq \mathsf{nt}(t_{2n}) \trianglerighteq t_{3n}))_{n\in\mathbb{N}}\,.$$

The *projective limit model* $\boldsymbol{\mathcal{I}}^{\infty}(\mathrm{prTA}_{\mathrm{si}})$ of $\mathrm{prTA}_{\mathrm{si}}$ is $\boldsymbol{\mathcal{I}}^{\infty}(\mathrm{prBTA}_{\mathrm{nt}})$ expanded with the operations defined by

$$\begin{aligned}
\|^n_{h,s}((t_{1n})_{n\in\mathbb{N}},\ldots,(t_{mn})_{n\in\mathbb{N}}) &= (\pi_n(\|^n_{h,s}(t_{1n},\ldots,t_{mn})))_{n\in\mathbb{N}}\,,\\
\between^{n,i}_{h,s}((t_{1n})_{n\in\mathbb{N}},\ldots,(t_{mn})_{n\in\mathbb{N}}) &= (\pi_n(\between^{n,i}_{h,s}(t_{1n},\ldots,t_{mn})))_{n\in\mathbb{N}}\,,\\
\mathsf{S}_{\mathsf{D}}((t_n)_{n\in\mathbb{N}}) &= (\pi_n(\mathsf{S}_{\mathsf{D}}(t_n)))_{n\in\mathbb{N}}
\end{aligned}$$

as interpretations of the additional operators of $\mathrm{prTA}_{\mathrm{si}}$. On the right-hand side of these equations, the symbols $\trianglelefteq\,\mathsf{nt}(_)\,\trianglerighteq$, $\|^n_{h,s}$, $\between^{n,i}_{h,s}$, and S_{D} denote the interpretation of the operators $\trianglelefteq\,\mathsf{nt}(_)\,\trianglerighteq$, $\|^n_{h,s}$, $\between^{n,i}_{h,s}$, and S_{D} in $\boldsymbol{\mathcal{I}}(\mathrm{prTA}_{\mathrm{si}})$.

It is straightforward to check that the operations with which $\boldsymbol{\mathcal{I}}^{\infty}(\mathrm{prBTA}_{\mathrm{nt}})$ is expanded always yield elements of $\boldsymbol{\mathcal{T}}^{\infty}(\mathrm{prBTA}_{\mathrm{nt}})$. It follows immediately from the construction of $\boldsymbol{\mathcal{I}}^{\infty}(\mathrm{prTA}_{\mathrm{si}})$ and Theorem 2 that $\boldsymbol{\mathcal{I}}^{\infty}(\mathrm{prTA}_{\mathrm{si}})$ is really a projective limit model of $\mathrm{prTA}_{\mathrm{si}}$.

References

1. Andova, S., Georgievska, S.: On compositionality, efficiency, and applicability of abstraction in probabilistic systems. In: Nielsen, M., et al. (eds.) SOFSEM 2009. Lecture Notes in Computer Science, vol. 5404, pp. 67–78. Springer-Verlag (2009)
2. Baeten, J.C.M., Bergstra, J.A., Smolka, S.A.: Axiomatizing probabilistic processes: ACP with generative probabilities. Information and Computation 121(2), 234–255 (1995)
3. Baeten, J.C.M., Weijland, W.P.: Process Algebra, Cambridge Tracts in Theoretical Computer Science, vol. 18. Cambridge University Press, Cambridge (1990)
4. Bergstra, J.A., Bethke, I., Ponse, A.: Cancellation meadows: A generic basis theorem and some applications. Computer Journal 56(1), 3–14 (2013)
5. Bergstra, J.A., Hirshfeld, Y., Tucker, J.V.: Meadows and the equational specification of division. Theoretical Computer Science 410(12–13), 1261–1271 (2009)
6. Bergstra, J.A., Loots, M.E.: Program algebra for sequential code. Journal of Logic and Algebraic Programming 51(2), 125–156 (2002)
7. Bergstra, J.A., Middelburg, C.A.: Thread algebra with multi-level strategies. Fundamenta Informaticae 71(2–3), 153–182 (2006)
8. Bergstra, J.A., Middelburg, C.A.: Thread algebra for strategic interleaving. Formal Aspects of Computing 19(4), 445–474 (2007)
9. Bergstra, J.A., Middelburg, C.A.: A thread algebra with multi-level strategic interleaving. Theory of Computing Systems 41(1), 3–32 (2007)
10. Bergstra, J.A., Middelburg, C.A.: Distributed strategic interleaving with load balancing. Future Generation Computer Systems 24(6), 530–548 (2008)
11. Bergstra, J.A., Middelburg, C.A.: Program algebra with a jump-shift instruction. Journal of Applied Logic 6(4), 553–563 (2008)
12. Bergstra, J.A., Middelburg, C.A.: Instruction sequence notations with probabilistic instructions. `arXiv:0906.3083v1 [cs.PL]` (June 2009)
13. Bergstra, J.A., Middelburg, C.A.: A thread calculus with molecular dynamics. Information and Computation 208(7), 817–844 (2010)
14. Bergstra, J.A., Middelburg, C.A.: Instruction sequence processing operators. Acta Informatica 49(3), 139–172 (2012)

15. Bergstra, J.A., Middelburg, C.A.: Instruction Sequences for Computer Science, Atlantis Studies in Computing, vol. 2. Atlantis Press, Amsterdam (2012)
16. Bergstra, J.A., Ponse, A.: Combining programs and state machines. Journal of Logic and Algebraic Programming 51(2), 175–192 (2002)
17. Bergstra, J.A., Ponse, A.: Signed meadow valued probability mass functions. `arXiv:1307.5173v1 [math.LO]` (July 2013)
18. Bergstra, J.A., Tucker, J.V.: The rational numbers as an abstract data type. Journal of the ACM 54(2), Article 7 (2007)
19. Gay, S.J.: Quantum programming languages: Survey and bibliography. Mathematical Structures in Computer Science 16(4), 581–600 (2006)
20. van Glabbeek, R.J., Smolka, S.A., Steffen, B.: Reactive, generative and stratified models of probabilistic processes. Information and Computation 121(1), 59–80 (1995)
21. Gosling, J., Joy, B., Steele, G., Bracha, G.: The Java Language Specification. Addison-Wesley, Reading, MA, second edn. (2000)
22. He Jifeng, Seidel, K., McIver, A.K.: Probabilistic models for the guarded command language. Science of Computer Programming 28(2–3), 171–192 (1997)
23. Hejlsberg, A., Wiltamuth, S., Golde, P.: C# Language Specification. Addison-Wesley, Reading, MA (2003)
24. Hoare, C.A.R.: Communicating Sequential Processes. Prentice-Hall, Englewood Cliffs (1985)
25. Hodges, W.A.: Model Theory, Encyclopedia of Mathematics and Its Applications, vol. 42. Cambridge University Press, Cambridge (1993)
26. Itai, A., Rodeh, M.: Symmetry breaking in distributed networks. Information and Computation 88(1), 60–87 (1990)
27. Milner, R.: Communication and Concurrency. Prentice-Hall, Englewood Cliffs (1989)
28. Rabin, M.O.: Probabilistic algorithms. In: Traub, J.F. (ed.) Algorithms and Complexity: New Directions and Recent Results, pp. 21–39. Academic Press, New York (1976)
29. Rabin, M.O.: Probabilistic algorithm for testing primality. Journal of Number Theory 12(1), 128–138 (1980)
30. Sabelfeld, A., Sands, D.: Probabilistic noninterference for multi-threaded programs. In: Computer Security Foundations Workshop 2000. pp. 200–214. IEEE Computer Society Press (2000)
31. Sannella, D., Tarlecki, A.: Algebraic preliminaries. In: Astesiano, E., Kreowski, H.J., Krieg-Brückner, B. (eds.) Algebraic Foundations of Systems Specification, pp. 13–30. Springer-Verlag, Berlin (1999)
32. Shor, P.W.: Algorithms for quantum computation: Discrete logarithms and factoring. In: FOCS '94. pp. 124–134. IEEE Computer Society Press (1994)
33. Waldspurger, C.A., Weihl, W.E.: Lottery scheduling: Flexible proportional-share resource management. In: OSDI '94. pp. 1–12. USENIX Association (1994)
34. Wirsing, M.: Algebraic specification. In: van Leeuwen, J. (ed.) Handbook of Theoretical Computer Science, vol. B, pp. 675–788. Elsevier, Amsterdam (1990)